CLAIMS

What is claimed is:

5   1.   A router comprising:

a computer readable medium;

a plurality of transmission links for receiving and sending data;

retrieving means for retrieving a destination address

10   from a data packet;

hashing means for hashing the destination address to determine a table index into a table; and

reading means for reading a target address from a table entry using the table index, wherein the target address has

15   been related to and stored in the table entry based on a computed value from a relation computation using the table index and the target address as operands in the relation computation; and

modifying means for modifying the data packet by

20   storing the target address in the data packet as a next-hop destination address.

2.   The router of claim 1 further comprising relating means for relating a particular table entry to a target address

25   wherein:

generating means for generating, for each target address in the set of target addresses, a computed value using the table index for the particular table entry and a target address as operands in the relation computation to

30   obtain a set of computed values;

AUS9-2000-0336-US1

       choosing means for choosing a computed value from the set of computed values based upon a mathematical relationship among the set of computed values; and

       determining means for determining a related target

5   address for the particular table entry based on the chosen computed value, wherein the chosen computed value was computed using the related target address as input.

     3.   The router of claim 2 further comprising:

10   obtaining means for obtaining a set of target addresses; and

       relating means for relating, for each table entry, a target address from the set of target addresses to a table entry such that each table entry is related with only one

15   target address; and

       storing means for storing in each table entry its related target address.

     4.   A routing method in a data processing system comprising

20   the steps of:

       receiving a data packet;

       retrieving a destination address from the data packet;

       hashing the destination address to determine a table index into a table in a computer readable medium; and

25   reading a target address from a table entry using the table index, wherein the target address has been related to and stored in the table entry based on a computed value from a relation computation using the table index and the target address as operands in the relation computation;

30   modifying the data packet by storing the target address in the data packet as a next-hop destination address; and

AUS9-2000-0336-US1

transmitting the modified data packet.

5. The method of claim 4 further comprising a step of relating a particular table entry to a target address in

5   which:

for each target address in the set of target addresses, generating a computed value using the table index for the particular entry and a target address as operands in the relation computation to obtain a set of computed values;

10   choosing a computed value from the set of computed values based upon a mathematical relationship among the set of computed values; and

determining a related target address for the particular entry based on the chosen computed value, wherein the chosen

15   computed value was computed using the related target address as input.

6. The method of claim 5 further comprising:

obtaining a set of target addresses; and

20   for each table entry, relating a target address from the set of target addresses to a table entry such that each table entry is related with only one target address; and

for each table entry, storing in a table entry its related target address.

25

7. A method in a data processing system for mapping a source identifier to a target identifier in a set of target identifiers, the method comprising the steps of:

managing a data structure in a computer readable

30   medium, wherein each entry within the data structure stores information associated with a single target identifier and

AUS9-2000-0336-US1

wherein a single target identifier is related to one or more entry locations;

hashing the source identifier to a location identifier of an entry in the data structure; and

5   retrieving information associated with the target identifier from the entry in the data structure using the location identifier; and

obtaining a mapped target identifier from the retrieved information associated with the target identifier,

10   wherein a processing speed with which the source identifier is mapped to the mapped target identifier is independent of a total number of target identifiers in the set of target identifiers.

15   8.   The method of claim 7 wherein the method for mapping the source identifier to the target identifier is stable with respect to changes in the set of target identifiers.

9.   A method in a data processing system for mapping a
20   source identifier to a target identifier, the method comprising the steps of:

hashing the source identifier to determine a table index into a table in a computer readable medium; and

reading the target identifier from a table entry using
25   the table index, wherein the target identifier has been related to and stored in the table entry based on a computed value from a relation computation using the table index and the target identifier as operands in the relation computation.

30

AUS9-2000-0336-US1

10. The method of claim 9 further comprising a step of relating a particular table entry to a target identifier in which:

for each target identifier in the set of target

5   identifiers, generating a computed value using the table index for the particular table entry and a target identifier as operands in the relation computation to obtain a set of computed values;

choosing a computed value from the set of computed

10   values based upon a mathematical relationship among the set of computed values; and

determining a related target identifier for the particular entry based on the chosen computed value, wherein the chosen computed value was computed using the related

15   target identifier as input.

11. The method of claim 10 further comprising, prior to the step of reading the target identifier from the table entry:

obtaining a set of target identifiers; and

20   for each table entry, relating a target identifier from the set of target identifiers to a table entry such that each table entry is related with only one target identifier; and

for each table entry, storing in a table entry its

25   related target identifier.

12. The method of claim 10 further comprising:

dynamically removing a target identifier from a set of target identifiers to obtain a modified set of target

30   identifiers;

AUS9-2000-0336-US1

for each table entry previously related to the removed target identifier, newly relating a target identifier from the modified set of target identifiers to a table entry such that each table entry is related with only one target

5   identifier; and

for each table entry previously related to the removed target identifier, storing in a table entry its newly related target identifier.

10   13.   The method of claim 10 further comprising:

dynamically adding a target identifier to a set of target identifiers to obtain a modified set of target identifiers;

for each table entry, relating a target identifier from

15   the modified set of target identifiers to a table entry such that each table entry is related with only one target identifier; and

for each table entry, storing in a table entry its related target identifier if its related target identifier

20   differs from a target identifier previously stored in the table entry.

14.   The method of claim 9 wherein the relation computation further comprises:

25   receiving the table index and the target identifier as operands for the relation computation;

hashing the table index to generate a first hash value;

hashing the target identifier to generate a second hash value; and

30   hashing the first hash value and the second hash value to generate a computed value.

AUS9-2000-0336-US1

15. The method of claim 9 further comprising:

obtaining a set of target identifiers, wherein each target identifier identifies a computational resource such

5   that each target identifier is related with only one computational resource.

16. The method of claim 9 further comprising:

associating a computational resource with a subset of a

10   set of target identifiers, wherein each target identifier in the set of target identifiers is related with only one computational resource, wherein each target identifier in the subset of target identifiers identifies the computational resource, and wherein a size of the subset of

15   target identifiers is proportional to a computational capacity of the computational resource.

17. The method of claim 9 wherein the source identifier is a network protocol address.

20

18. The method of claim 9 wherein the target identifier is a network physical address.

19. The method of claim 9 wherein the target identifier is

25   a Uniform Resource Identifier (URI).

20. An apparatus for mapping a source identifier to a target identifier comprising:

first hashing means for hashing the source identifier

30   to determine a table index into a table in a computer readable medium; and

AUS9-2000-0336-US1

reading means for reading the target identifier from a table entry using the table index, wherein the target identifier has been related to and stored in the table entry based on a computed value from a relation computation using

5   the table index and the target identifier as operands in the relation computation.

21.   The apparatus of claim 20 further comprising first relating means for relating a particular table entry to a

10   target identifier wherein:

generating means for generating, for each target identifier in the set of target identifiers, a computed value using the table index for the particular entry and a target identifier as operands in the computation to obtain a

15   set of computed values;

choosing means for choosing a computed value from the set of computed values based upon a mathematical relationship among the set of computed values; and

determining means for determining a related target

20   identifier for the particular table entry based on the chosen computed value, wherein the chosen computed value was computed using the related target identifier as input.

22.   The apparatus of claim 21 further comprising:

25   first obtaining means for obtaining a set of target identifiers; and

second relating means for relating, for each table entry, a target identifier from the set of target identifiers to a table entry such that each table entry is

30   related with only one target identifier; and

AUS9-2000-0336-US1

first storing means for storing in each table entry its related target identifier.

23. The apparatus of claim 21 further comprising:

5   removing means for dynamically removing a target identifier from a set of target identifiers to obtain a modified set of target identifiers;

third relating means for newly relating, for each table entry previously related to the removed target identifier, a

10   target identifier from the modified set of target identifiers to a table entry such that each table entry is related with only one target identifier; and

second storing means for storing, in each table entry previously related to the removed target identifier, its

15   newly related target identifier.

24. The apparatus of claim 21 further comprising:

adding means for dynamically adding a target identifier to a set of target identifiers to obtain a modified set of

20   target identifiers; and

fourth relating means for relating, for each table entry, a target identifier from the modified set of target identifiers to a table entry such that each table entry is related with only one target identifier; and

25   third storing means for storing in each table entry its related target identifier if its related target identifier differs from a target identifier previously stored in the table entry.

AUS9-2000-0336-US1

25. The apparatus of claim 20 further comprising a relation
computation means for computing the relation computation
wherein:

5   receiving means for receiving the table index and the
target identifier as operands for the relation computation;

         second hashing means for hashing the table index to
generate a first hash value;

         third hashing means for hashing the target identifier

10 to generate a second hash value; and

         fourth hashing means for hashing the first hash value
and the second hash value to generate a computed value.

26. The apparatus of claim 20 further comprising:

15   second obtaining means for obtaining a set of target
identifiers, wherein each target identifier identifies a
computational resource such that each target identifier is
related with only one computational resource.

20 27. The apparatus of claim 20 further comprising:
         associating means for associating a computational
resource with a subset of a set of target identifiers,
wherein each target identifier in the set of target
identifiers is related with only one computational resource,

25 wherein each target identifier in the subset of target
identifiers identifies the computational resource, and
wherein a size of the subset of target identifiers is
proportional to a computational capacity of the
computational resource.

30

28. The apparatus of claim 20 wherein the source identifier is a network protocol address.

29. The apparatus of claim 20 wherein the target identifier is a network physical address.

30. The apparatus of claim 20 wherein the target identifier is a Uniform Resource Identifier (URI).

31. A computer program product on a computer readable medium for use in a data processing system for mapping a source identifier to a target identifier, the computer program product comprising:

instructions for hashing the source identifier to determine a table index into a table in a computer readable medium; and

instructions for reading the target identifier from a table entry using the table index, wherein the target identifier has been related to and stored in the table entry based on a computed value from a relation computation using the table index and the target identifier as operands in the relation computation.

32. The computer program product of claim 31 further comprising instructions for relating a particular table entry to a target identifier including:

instructions for generating, for each target identifier in the set of target identifiers, a computed value using the table index for the particular table entry and a target identifier as operands in the relation computation to obtain a set of computed values;

AUS9-2000-0336-US1

instructions for choosing a computed value from the set of computed values based upon a mathematical relationship among the set of computed values; and

instructions for determining a related target
5   identifier for the particular entry based on the chosen computed value, wherein the chosen computed value was computed using the related target identifier as input.

33. The computer program product of claim 32 further
10   comprising:

instructions for obtaining a set of target identifiers; and

instructions for relating, for each table entry, a target identifier from the set of target identifiers to a
15   table entry such that each table entry is related with only one target identifier; and

instructions for storing in each table entry its related target identifier.

20   34. The computer program product of claim 32 further comprising:

instructions for dynamically removing a target identifier from a set of target identifiers to obtain a modified set of target identifiers;
25   instructions for newly relating, for each table entry previously related to the removed target identifier, a target identifier from the modified set of target identifiers to a table entry such that each table entry is related with only one target identifier; and

AUS9-2000-0336-US1

instructions for storing, in each table entry previously related to the removed target identifier, its newly related target identifier.

5   35.   The computer program product of claim 32 further comprising:

instructions for dynamically adding a target identifier to a set of target identifiers to obtain a modified set of target identifiers;

10   instructions for relating, for each table entry, a target identifier from the modified set of target identifiers to a table entry such that each table entry is related with only one target identifier; and

instructions for storing in each table entry its
15   related target identifier if its related target identifier differs from a target identifier previously stored in the table entry.

36.   The computer program product of claim 31 wherein the
20   relation computation further comprises:

instructions for receiving the table index and the target identifier as operands for the relation computation;

instructions for hashing the table index to generate a first hash value;

25   instructions for hashing the target identifier to generate a second hash value; and

instructions for hashing the first hash value and the second hash value to generate a computed value.

AUS9-2000-0336-US1

37. The computer program product of claim 31 further comprising:

instructions for obtaining a set of target identifiers,

5 wherein each target identifier identifies a computational resource such that each target identifier is related with only one computational resource.

38. The computer program product of claim 31 further

10 comprising:

instructions for associating a computational resource with a subset of a set of target identifiers, wherein each target identifier in the set of target identifiers is related with only one computational resource, wherein each

15 target identifier in the subset of target identifiers identifies the computational resource, and wherein a size of the subset of target identifiers is proportional to a computational capacity of the computational resource.

20 39. The computer program product of claim 31 wherein the source identifier is a network protocol address.

40. The computer program product of claim 31 wherein the target identifier is a network physical address.

25

41. The computer program product of claim 31 wherein the target identifier is a Uniform Resource Identifier (URI).

AUS9-2000-0336-US1

42. A method in a data processing system for mapping a
source identifier to a target identifier, the method
comprising the steps of:

5   hashing the source identifier to determine a location
identifier of an entry in a data structure in a computer
readable medium; and

        reading information associated with the target
identifier from the entry in the data structure using the
10   location identifier, wherein the information associated with
the target identifier has been related to and stored in the
entry based on a computed value from a relation computation
using the location identifier and the target identifier as
operands in the relation computation.

15

43. The method of claim 42 further comprising a step of
relating a particular entry in the data structure to a
target identifier in which:

        for each target identifier in the set of target
20   identifiers, generating a computed value using the location
identifier of the particular entry in the data structure and
a target identifier as operands in the relation computation
to obtain a set of computed values;

        choosing a computed value from the set of computed
25   values based upon a mathematical relationship among the set
of computed values; and

        determining a related target identifier for the
particular entry based on the chosen computed value, wherein
the chosen computed value was computed using the related
30   target identifier as input.

AUS9-2000-0336-US1

44. The method of claim 43 further comprising, prior to the step of reading the information associated with the target identifier from the entry in the data structure:

obtaining a set of target identifiers; and

5   for each entry in the data structure, relating a target identifier from the set of target identifiers to an entry in the data structure such that each entry in the data structure is related with only one target identifier; and

for each entry in the data structure, storing in an
10   entry information associated with its related target identifier.
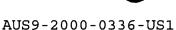
45. The method of claim 43 further comprising:

dynamically removing a target identifier from a set of
15   target identifiers to obtain a modified set of target identifiers;

for each entry in the data structure previously related to the removed target identifier, newly relating a target identifier from the modified set of target identifiers to an
20   entry in the data structure such that each entry in the data structure is related with only one target identifier;

for each entry in the data structure previously related to the removed target identifier, storing in an entry information associated with its newly related target
25   identifier.

46. The method of claim 43 further comprising:

dynamically adding a target identifier to a set of target identifiers to obtain a modified set of target
30   identifiers; and

AUS9-2000-0336-US1

     for each entry in the data structure, relating a target identifier from the modified set of target identifiers to an entry in the data structure such that each entry in the data structure is related with only one target identifier; and

5   for each entry in the data structure, storing in an entry information associated with its related target identifier if its related target identifier differs from a target identifier previously related to the entry in the data structure.

10

47. The method of claim 42 wherein the relation computation further comprises:

     receiving the location identifier and the target identifier as operands for the relation computation;

15   hashing the location identifier to generate a first hash value;

     hashing the target identifier to generate a second hash value; and

     hashing the first hash value and the second hash value

20 to generate a computed value.

48. The method of claim 42 further comprising:

     obtaining a set of target identifiers, wherein each target identifier identifies a computational resource such

25 that each target identifier is related with only one computational resource.

49. The method of claim 42 further comprising:

     associating a computational resource with a subset of a

30 set of target identifiers, wherein each target identifier in the set of target identifiers is related with only one

AUS9-2000-0336-US1

computational resource, wherein each target identifier in the subset of target identifiers identifies the computational resource, and wherein a size of the subset of target identifiers is proportional to a computational

5   capacity of the computational resource.

50.   The method of claim 42 further comprising:
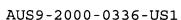   retrieving the target identifier using the information associated with the target identifier; and

10   performing a computational process on a computational resource identified by the target identifier.

51.   The method of claim 50 wherein the computational resource identified by the target identifier is a memory

15   resource.

52.   The method of claim 50 wherein the computational resource identified by the target identifier is a data processing system.

20

53.   The method of claim 42 wherein the information associated with the target identifier comprises the target identifier.

25   54.   The method of claim 42 wherein the data structure is a table, and the location identifier is a table index.

55.   The method of claim 42 wherein the source identifier is a network protocol address.

30

AUS9-2000-0336-US1

56. The method of claim 42 wherein the target identifier is a network physical address.

57. The method of claim 42 wherein the target identifier is
5   a Uniform Resource Identifier (URI).

58. An apparatus for mapping a source identifier to a target identifier comprising:
    first hashing means for hashing the source identifier
10 to determine a location identifier of an entry in a data structure in a computer readable medium; and
    reading means for reading information associated with the target identifier from the entry in the data structure using the location identifier, wherein the information
15 associated with the target identifier has been related to and stored in the entry based on a computed value from a relation computation using the location identifier and the target identifier as operands in the relation computation.

20 59. The apparatus of claim 58 further comprising first relating means for relating a particular entry in the data structure to a target identifier wherein:
    generating means for generating, for each target identifier in the set of target identifiers, a computed
25 value using the location identifier of the particular entry in the data structure and a target identifier as operands in the relation computation to obtain a set of computed values;
    choosing means for choosing a computed value from the set of computed values based upon a mathematical
30 relationship among the set of computed values; and

AUS9-2000-0336-US1

determining means for determining a related target
identifier for the particular entry based on the chosen
computed value, wherein the chosen computed value was
computed using the related target identifier as input.

5

60. The apparatus of claim 59 further comprising, prior to
the step of reading the information associated with the
target identifier from the entry in the data structure:

first obtaining means for obtaining a set of target
10 identifiers; and

second relating means for relating, for each entry in
the data structure, a target identifier from the set of
target identifiers to an entry in the data structure such
that each entry in the data structure is related with only
15 one target identifier; and

first storing means for storing, in each entry in the
data structure, information associated with its related
target identifier.

20 61. The apparatus of claim 59 further comprising:
removing means for dynamically removing a target
identifier from a set of target identifiers to obtain a
modified set of target identifiers;

third relating means for newly relating, for each entry
25 in the data structure previously related to the removed
target identifier, a target identifier from the modified set
of target identifiers to an entry in the data structure such
that each entry in the data structure is related with only
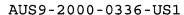one target identifier;

30 second storing means for storing, in each entry in the
data structure previously related to the removed target

identifier, information associated with its newly related target identifier.

62.   The apparatus of claim 59 further comprising:

5   adding means for dynamically adding a target identifier to a set of target identifiers to obtain a modified set of target identifiers; and

          fourth relating means for relating, for each entry in the data structure, a target identifier from the modified

10   set of target identifiers to an entry in the data structure such that each entry in the data structure is related with only one target identifier; and

          third storing means for storing, in each entry in the data structure, information associated with its related

15   target identifier if its related target identifier differs from a target identifier previously related to the entry in the data structure.

63.   The apparatus of claim 58 wherein the relation

20   computation further comprises:

          receiving means for receiving the location identifier and the target identifier as operands for the relation computation;

          second hashing means for hashing the location

25   identifier to generate a first hash value;

          third hashing means for hashing the target identifier to generate a second hash value; and

          fourth hashing means for hashing the first hash value and the second hash value to generate a computed value.
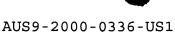
30

AUS9-2000-0336-US1

64.   The apparatus of claim 58 further comprising:

      second obtaining a set of target identifiers, wherein
each target identifier identifies a computational resource
such that each target identifier is related with only one
5   computational resource.

65.   The apparatus of claim 58 further comprising:

      associating means for associating a computational
resource with a subset of a set of target identifiers,
10   wherein each target identifier in the set of target
identifiers is related with only one computational resource,
wherein each target identifier in the subset of target
identifiers identifies the computational resource, and
wherein a size of the subset of target identifiers is
15   proportional to a computational capacity of the
computational resource.

66.   The apparatus of claim 58 further comprising:

      retrieving means for retrieving the target identifier
20   using the information associated with the target identifier;
and

      performing means for performing a computational process
on a computational resource identified by the target
identifier.

25

67.   The apparatus of claim 64 wherein the computational
resource identified by the target identifier is a memory
resource.

AUS9-2000-0336-US1

68. The apparatus of claim 64 wherein the computational resource identified by the target identifier is a data processing system.

5   69. The apparatus of claim 58 wherein the information associated with the target identifier comprises the target identifier.

70. The apparatus of claim 58 wherein the data structure
10 is a table, and the location identifier is a table index.

71. The apparatus of claim 58 wherein the source identifier is a network protocol address.

15 72. The apparatus of claim 58 wherein the target identifier is a network physical address.

73. The apparatus of claim 58 wherein the target identifier is a Uniform Resource Identifier (URI).
20

74. A computer program product on a computer readable medium for use in a data processing system for mapping a source identifier to a target identifier, the computer program product comprising:

25   instructions for hashing the source identifier to determine a location identifier of an entry in a data structure in a computer readable medium; and

        instructions for reading information associated with the target identifier from the entry in the data structure
30 using the location identifier, wherein the information associated with the target identifier has been related to

AUS9-2000-0336-US1

and stored in the entry based on a computed value from a relation computation using the location identifier and the target identifier as operands in the relation computation.

5   75.   The computer program product of claim 74 further comprising instructions for relating a particular entry in the data structure to a target identifier including:

instructions for generating, for each target identifier in the set of target identifiers, a computed value using the

10   location identifier of the particular entry in the data structure and a target identifier as operands in the relation computation to obtain a set of computed values;

instructions for choosing a computed value from the set of computed values based upon a mathematical relationship

15   among the set of computed values; and

instructions for determining a related target identifier for the particular entry based on the chosen computed value, wherein the chosen computed value was computed using the related target identifier as input.

20

76.   The computer program product of claim 75 further comprising:
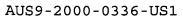
instructions for obtaining a set of target identifiers; and

25   instructions for relating, for each entry in the data structure, a target identifier from the set of target identifiers to an entry in the data structure such that each entry in the data structure is related with only one target identifier; and

instructions for storing, in each entry in the data structure, information associated with its related target identifier.

5   77.   The computer program product of claim 75 further comprising:

instructions for dynamically removing a target identifier from a set of target identifiers to obtain a modified set of target identifiers;

10   instructions for newly relating, for each entry in the data structure previously related to the removed target identifier, a target identifier from the modified set of target identifiers to an entry in the data structure such that each entry in the data structure is related with only

15   one target identifier;

instructions for storing, in each entry in the data structure previously related to the removed target identifier, information associated with its newly related target identifier.

20

78.   The computer program product of claim 75 further comprising:

instructions for dynamically adding a target identifier to a set of target identifiers to obtain a modified set of

25   target identifiers; and

instructions for relating, for each entry in the data structure, a target identifier from the modified set of target identifiers to an entry in the data structure such that each entry in the data structure is related with only
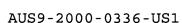
30   one target identifier; and

AUS9-2000-0336-US1

instructions for storing, in each entry in the data
structure, information associated with its related target
identifier if its related target identifier differs from a
target identifier previously related to the entry in the

5   data structure.

79.   The computer program product of claim 74 wherein the
relation computation further comprises:

instructions for receiving the location identifier and

10   the target identifier as operands for the relation
computation;

instructions for hashing the location identifier to
generate a first hash value;

instructions for hashing the target identifier to

15   generate a second hash value; and

instructions for hashing the first hash value and the
second hash value to generate a computed value.

80.   The computer program product of claim 74 further

20   comprising:

instructions for obtaining a set of target identifiers,
wherein each target identifier identifies a computational
resource such that each target identifier is related with
only one computational resource.

25

81.   The computer program product of claim 74 further
comprising:

instructions for associating a computational resource
with a subset of a set of target identifiers, wherein each

30   target identifier in the set of target identifiers is
related with only one computational resource, wherein each

AUS9-2000-0336-US1

target identifier in the subset of target identifiers
identifies the computational resource, and wherein a size of
the subset of target identifiers is proportional to a
computational capacity of the computational resource.

5

82. The computer program product of claim 74 further
comprising:

        instructions for retrieving the target identifier using
the information associated with the target identifier; and
10   instructions for performing a computational process on
a computational resource identified by the target
identifier.

83. The computer program product of claim 80 wherein the
15   computational resource identified by the target identifier
is a memory resource.

84. The computer program product of claim 80 wherein the
computational resource identified by the target identifier
20   is a data processing system.

85. The computer program product of claim 74 wherein the
information associated with the target identifier comprises
the target identifier.

25

86.   The computer program product of claim 74 wherein the
data structure is a table, and the location identifier is a
table index.

30   87. The computer program product of claim 74 wherein the
source identifier is a network protocol address.

AUS9-2000-0336-US1

88. The computer program product of claim 74 wherein the target identifier is a network physical address.

5 89. The computer program product of claim 74 wherein the target identifier is a Uniform Resource Identifier (URI).